Observational Studies

Introduction to Quantitative Social Science

Kosuke Imai
Harvard University / University of Tokyo

Summer 2022

Review of Randomized Control Trials

- Fundamental problem of causal inference:
  - Comparison between factual and counterfactual
  - Counterfactuals are not observed
- Solution: Randomized controlled trials (RCTs)
  - Treatment and control groups *identical on average*
  - Similar in all (observed and unobserved) characteristics
- Difference in average outcome between the two groups as an estimate of Sample Average Treatment Effect (SATE):

$$
\begin{aligned}
\text{difference-in-means estimator} &= \frac{1}{n_1}\sum_{i=1}^{n} T_i Y_i - \frac{1}{n-n_1}\sum_{i=1}^{n}(1-T_i)Y_i \\
\text{SATE} &= \frac{1}{n}\sum_{i=1}^{n}\{Y_i(1)-Y_i(0)\}
\end{aligned}
$$

- Examples of RCTs:
  - Causal effect of race on employment prospect
  - Causal effect of naming-and-shaming on turnout

Observational Studies

- Often, we can't randomize treatment for ethical and logistical reasons: e.g., smoking and lung cancer
- Observational studies: naturally assigned treatment

- Better external validity for generalization beyond experiment
- Weaker internal validity:
    - pre-treatment variables may differ between treatment and control groups
    - confounding bias due to these differences
    - selection bias from self-selection into treatment
    - statistical control needed
    - unobserved confounding poses a threat

Confounding

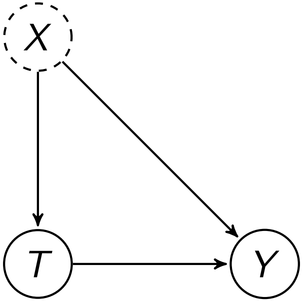

- Key assumption "Unconfoundedness": treatment and control groups comparable with respect to everything other than treatment
- How can we find a good comparison group?

Minimum Wage and Unemployment

- How does the increase in minimum wage affect employment?
- Current debate: federal minimum wage increase
- Many economists believe the effect is negative
  - especially for the poor
  - also for the whole economy
- Hard to randomize the minimum wage increase
- Two social scientists tested this using fast food chains in NJ and PA
- In 1992, NJ minimum wage increased from $4.25 to $5.05
  - Neighboring PA stays at $4.25
  - Observe employment in both states before and after increase
- NJ and (eastern) PA are similar
- Fast food chains in NJ and PA are similar: price, wages, products, etc.
- They are most likely to be affected by this increase

Name	Description
chain	name of fastfood restaurant chain
location	location of restaurants (centralNJ, northNJ, PA, shoreNJ, southNJ)
wageBefore	wage before the minimum wage increase
wageAfter	wage after the minimum wage increase
fullBefore	number of fulltime employees before the minimum wage increase
fullAfter	number of fulltime employees before the minimum wage increase
partBefore	number of parttime employees before the minimum wage increase
partAfter	number of parttime employees before the minimum wage increase

```
minwage <- read.csv("data/minwage.csv")
dim(minwage)

## [1] 358   8
```

Did the Minimum Wage Law Affect the Wages in NJ?

- Subset the data into NJ and PA

```
minwageNJ <- subset(minwage, subset = (location != "PA"))
minwagePA <- subset(minwage, subset = (location == "PA"))
```

- Compute the proportion of restaurants whose wage is less than $5.05

```
mean(minwageNJ$wageBefore < 5.05) # NJ before
## [1] 0.911
mean(minwageNJ$wageAfter < 5.05)  # NJ after
## [1] 0.00344
mean(minwagePA$wageBefore < 5.05) # PA before
## [1] 0.94
mean(minwagePA$wageAfter < 5.05)  # PA after
## [1] 0.955
```

Are the NJ and PA Restaurants Comparable?

- Average wages before the increase of minimum wage:

```
mean(minwageNJ$wageBefore)
## [1] 4.61
```

```
mean(minwagePA$wageBefore)
## [1] 4.65
```

- Prior proportion of fulltime employment:

```
minwageNJ$fullPropBefore <- minwageNJ$fullBefore /
    (minwageNJ$fullBefore + minwageNJ$partBefore)
minwagePA$fullPropBefore <- minwagePA$fullBefore /
    (minwagePA$fullBefore + minwagePA$partBefore)
mean(minwageNJ$fullPropBefore)
## [1] 0.297
mean(minwagePA$fullPropBefore)
## [1] 0.31
```

Cross-section Comparison

- Compare NJ and PA using the data after the increase
- The treatment and control groups are assumed to be identical on average in terms of all confounders
- What confounders are missing from the data?
- Compute the proportion of fulltime employees after the increase:

```
minwageNJ$fullPropAfter <- minwageNJ$fullAfter /
(minwageNJ$fullAfter + minwageNJ$partAfter)
minwagePA$fullPropAfter <- minwagePA$fullAfter /
(minwagePA$fullAfter + minwagePA$partAfter)
```

- The estimated SATE:

```
mean(minwageNJ$fullPropAfter) -
  mean(minwagePA$fullPropAfter)
## [1] 0.0481
```

Before-and-after Comparison

- State-specific confounders for cross-section comparison
- Compare NJ before and after
- What might be time-varying confounders?

```
NJdiff <- mean(minwageNJ$fullPropAfter) -
    mean(minwageNJ$fullPropBefore)
NJdiff
## [1] 0.0239
```

Difference-in-Differences

- Key Idea: use PA before-and-after difference to figure out what would have happened in NJ without the increase
- NJ before-and-after difference addresses within-state confounding

- Parallel time trend assumption
- Estimate the sample average treatment effect for the treated (SATT)
- The SATE of minimum wage increase in NJ

```
PAdiff <- mean(minwagePA$fullPropAfter) -
  mean(minwagePA$fullPropBefore)
NJdiff - PAdiff
## [1] 0.0616
```

Visualizing Difference-in-Differences

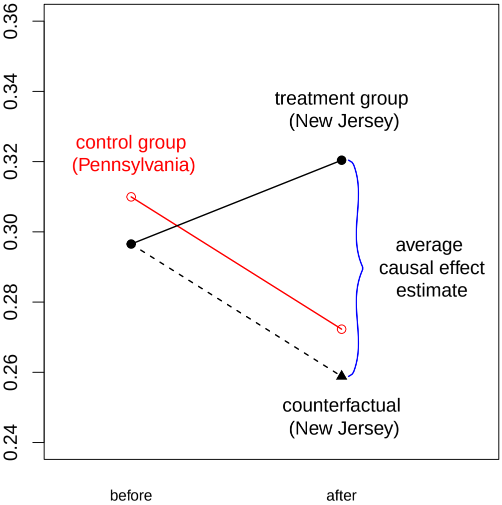

Summary of 3 Identification Strategies

1. Cross-section comparison
   - Compare treated units with control units after the treatment
   - Assumption: the treated and control units are comparable
   - Possible unit-specific confounding
2. Before-and-after comparison
   - Compare the same units before and after the treatment
   - Assumption: no time-varying confounding
3. Difference-in-Differences
   - Assumption: parallel time trend
   - Under this assumption, it accounts for both unit-specific and time-varying confounding

Neither approach is best. They require different assumptions.